An Analysis of Technological and Marketplace Developments from 2003 to Present and Their Impact on the Effectiveness of CAN-SPAM

Paul Q. Judge
CipherTrust, Inc.[a]
4800 North Point Pkwy
Alpharetta, GA 30022
paul.judge@ciphertrust.com

1 Introduction

Prior to the enactment of CAN-SPAM, spam grew quickly from a little known nuisance to a widespread problem that affected almost every user of the Internet. Efforts to solve the problem have focused on three fronts: technology, user awareness and legislation. While fundamentally related, each of these fronts has progressed over the years somewhat independently.

Dealing with spam requires a long-term cure as well as some immediate symptom relief. Currently, anti-spam products that are available for e-mail gateways and for desktops can detect and block well over 90 percent of spam.[1] These solutions are effective at relieving the spam symptoms for organizations that have deployed them.

Stopping spam globally requires removing the incentive to spam. The overall solution to spam thus must aim to reduce the profitability of sending unwanted e-mail. Just as in any other business, the profit in spamming is equal to revenues minus expenses. In spamming, expenses include the cost of obtaining the lists of e-mail addresses and the cost of sending the messages. Revenue is equal to the number of spam messages actually received by the intended recipients' multiplied by the response rate multiplied by the profit per item. Both expenses and revenue can be affected by user education, which

[a] Affiliation included for identification purposes only. The opinions expressed in this paper are those of the author alone, and not necessarily the opinions of any other person, entity, agency, or organization.

influences the recipients' response rate, as well as the spammer's difficulty and costs involved in obtaining e-mail addresses. In addition, the difference between the number of spam messages sent and the number received is a product of the effectiveness and deployment rate of anti-spam technologies.

Anti-spam legislation provides strong deterrence. The relationship between technology and legislation is a familiar one. For example, consider the problem of computer intrusions: there is technology available—such as firewalls and intrusion-detection systems—to protect resources, and the technology is supported by legislation to allow prosecution of and litigation against those who are determined to circumvent the technology. This same relationship is seen in something as basic as property theft—burglar alarms and door locks provide protection, and they are supported by laws that provide strong deterrence.

In this report, we analyze the changes that have occurred since CAN-SPAM[2] was enacted to examine any impact on effectiveness of the provisions of the law. We offer taxonomy of anti-spam technologies and highlight the technological advances since the bill was enacted, and their affect on CAN-SPAM.

2 How Spammers Take Advantage of the Email System

Simple Mail Transfer Protocol (SMTP) is the system that defines how email works[3]. It explains how email travels from the sender to the sender's mail server to the recipient's mail server to the recipient. The format of the email messages is defined in another Internet standard document[4]. The channels used to transmit Internet applications are called ports; email is transmitted on port 25. In this section we highlight how the email format and transport protocol are relevant to the spam problem because spammers are leveraging weaknesses in these protocols to carry out their activities. Readers can reference[5] for an overview of the technical workings of email.

It is important to examine how the very properties of email that make it attractive are the properties that cause the problems of spam. These properties include:

1) Unbalanced Cost System: Email is one of the few communication forms in which the cost is borne by the recipient instead of the sender. For example, in postal mail the sender must pay for postage. In the phone system, the caller incurs the long distance charges. However, in email, the cost of sending an email is close to zero while the recipient incurs the cost of bandwidth, processing time, hardware costs, storage costs, mail server software licenses, and lost productivity in dealing with the messages. This makes it extremely efficient for a spammer to carry out large scale spam attacks.

2) Rapid Distribution Model: When normal users send email they have come to expect the message to show up in the recipient's inbox almost instantaneously. The elements of SMTP that allow that to happen also allow wide scale and rapid distribution of all types of messages in the email ecosystem. The scaling model of email allows an attacker to send a single message addressed to multiple recipients and have other mail servers do the work of splitting it into multiple messages and delivering each to the correct destination.

3) Open Infrastructure Access: Email separates the routing information from the identification information displayed to the user. Additionally, the email system has no native methods for authenticating the routing or purported identification. This means anyone can send an email message claiming to be from anyone else and there is no way of the recipient distinguishing. This disadvantages recipients, email filters, and law enforcement.

Today these three properties still exist and spammers are still being exploited by spammers. Anti-spam technologies today are focused on protecting users and the infrastructure given this environment. Some anti-spam approaches focus on minimizing the exposure of these three properties.

3 Anti-Spam Technologies

In this section we explore the landscape of anti-spam technologies. We provide taxonomy to show the relationship between different approaches. Over the last few years, there has been significant commercial activity in the anti-spam market[6]. There have been many new companies formed to offer solutions to the problem as well as interest from traditional computing companies. This competitiveness has fueled innovation and the

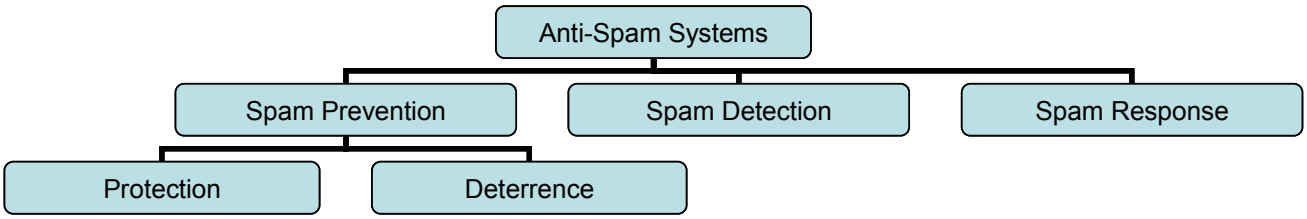

Figure 1: Types of Anti-Spam Systems

community has benefited from advances in anti-spam approaches over the last few years. One view of this progress can be seen by comparing the effectiveness rates of anti-spam solutions published in product reviews over the last few years. In early 2003 the leading six solutions averaged 84% spam catch rate[7] while in 2005 a similar review by the same publication the leading six solutions averaged a 98% spam catch rate[8]. Figure 1 shows the three types of anti-spam systems. These include the following: 1) spam prevention: These systems include protection and deterrence systems that focus on keeping spam from happening; 2) spam detection: These systems aim to identify incoming spam in order to take some action on it; 3) spam response: These systems are responsible for taking action once spam messages or spam sources are identified. These three types of systems are closely related. For example, detection systems leverage information from prevention systems and response systems act about the decisions made by detection systems.

In a typical security problem, the order of solutions is protection and deterrence to keep the threat from occurring and detection is relied upon as a last resort. Early work in anti-spam approached the problem from the opposite direction. Most anti-spam systems that were in place in 2003 focused on spam detection. This made sense years ago when a

smaller proportion of email was spam. Today, as a larger proportion of email is spam it is even more challenging to identify the spam messages than to just focus on finding the non-spam messages. Spam comes from a large number of rapidly changing machines around the world that are attempting to mask their identity and using other tricks to make them difficult to detect. Good mail, on the other hand, comes from a relatively small set of more static senders. This is why over the last several years much of the industry's efforts have gone into systems that also heavily consider good mail as part of the overall anti-spam system.

3.1 Spam Detection Systems

Spam detection systems attempt to distinguish spam messages from non-spam messages. The task is challenging because it consists of a machine making a decision for a human. This decision must typically be made in near real-time for hundreds of thousands messages going to thousands of individuals. Spam detection can take place in several places in the email lifecycle: 1) outbound mail traffic at an organization, 2) as an intermediary before the message reaches the recipient's network, 3) at the gateway of the recipient's network, 4) on the recipient's mail server or 5) at the recipient's desktop. The machine must make the spam/not spam decision based on a very limited amount of information. The only knowledge the machine has is based on four general questions about the message: (1) Who is it from? (2) What is in it? (3) How was it sent? (4) Where was it sent? Figure 2 shows the spectrum of approaches.

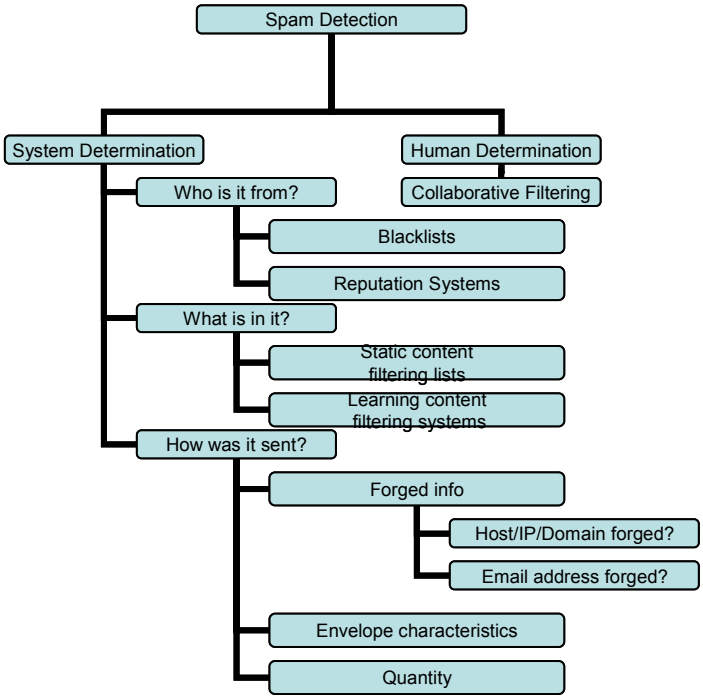

Figure 2: Types of Spam Detection Systems

Systems that identify spam based on the source began years ago with traditional blacklists. These are typically lists of IP addresses of computers that have been determined to send spam. Historically, there has been much controversy around blacklists because of the subjectivity of getting placed on the list and the difficulty of being removed. There are many public blacklists that exist with varying spam effectiveness and false positive rates. Whitelists have been developed that maintain lists of legitimate senders. A problem with whitelists and blacklists is they left a sizable set of senders in the middle of the spectrum that were not classified. This is because there was not enough credible information or feedback to make a binary decision about the sender. To address this problem, reputation systems for email became a topic of interest in the industry around the end of 2003 and beginning of 2004[9]. During this time, commercial reputation systems became available from several companies including BrightMail, CipherTrust, and Cloudmark.

Reputation systems focus on monitoring sender activity, analyzing the behavior, and determining a reputation for the sender. The reputation can be thought of like a credit score. In the credit scoring system, every individual is given a credit score based on some analysis of one's past financial behavior. In a similar nature, reputation systems aim to assign every computer on the Internet an email reputation. Typically, a reputation system uses the IP address of a computer as the identity and a reputation is built for that identity. There are many characteristics that can be analyzed to build the reputation including the mail volume, sending rate, complaint rate, and received messages. These characteristics are analyzed to determine patterns that are consistent with legitimate senders, illegitimate senders and those in between. The reputation is usually a positive or negative score across some defined range. Accreditation systems involve a third-party that monitors the sender's behavior, practices and standards. The system sometimes involves the sender placing a bond or payment to guarantee their good behavior. Critics of such systems suggest that they allow spammers to simply pay to obtain preferential treatment of their messages.[10] With the existence of better reputation systems, the manual processes involved in accreditation systems are not necessary.

Today, the leading commercial solutions utilize reputation systems. Some reputation systems are able to classify 70% of email traffic alone. Although the precision of a reputation system may provide a range of thousands of values, for a user there are four major categories that matter: unwanted, suspicious, unknown, and good. The real-time nature of these systems allows them to be effective in today's environment that involves hundreds of thousands of new zombie machines sending spam every day.

Content classification techniques examine the contents of the message to look for signs of spam-like content. There are two broad classes of techniques being used. The traditional approach uses described content-static lists of common spam words compiled manually in an ad hoc manner. At the end of 2002, a popular essay was written that introduced Bayesian filtering, a statistical and adaptive content analysis approach, to the broader technical community.[11] Through 2003 many open-source systems were released that built on this approach and several commercial solutions began to utilize it. The anti-spam community uses the word Bayesian filtering broadly to describe many other types of

statistical filtering. The two main contributions of these approaches is that the word lists that are used for classification are statistical lists based on actual probabilities of the words occurring in spam or non-spam mail as opposed to the manual ad hoc approach used previously. The other contribution is that these statistical approaches made it possible to have the machine adapt to new spam messages by retraining on new messages. These seemingly simple mechanisms provided better effectiveness rates with better accuracy than many commercial solutions did at the time. So within the year of 2003, the average catch rate across the industry improved as most utilized some form of adaptive statistical content filtering. Since 2003, there have been several advancements in statistical filtering approaches focused on improving efficiency and effectiveness as well as in being robust against spammer countermeasures.[12]

Now we briefly examine the other two of the four main questions asked by spam detection systems. The question 'How was it sent?' refers to systems that use rules or heuristics to examine the email transaction. These rules seek out spam-like behavior such as mismatched domain names, forged headers, or fake sender information. The question 'Where was it sent?' refers to the use of spamtraps, fake email addresses used to attract and collect spam. Messages received at these systems are used to train different types of spam detection.

3.2 Spam Prevention Systems

Spam prevention systems include deterrence and protection approaches as shown in Figure 3. Deterrence systems aim to discourage the act of spamming, whereas protection systems aim to shield systems from exposure to spam.

Protection systems take more of a fail-closed approach to the problem. That is they assume messages are spam until proven otherwise. In most deployments, these systems are used as a first layer before detection techniques so overall the approach is not a fail-closed system. A basic type of protection system is a whitelist that is a list of desired senders. This list can be made of email addresses, domain names, or IP addresses. Better implementations verify the whitelists using authentication techniques, tokens or

disposable email addresses. In 2003 challenge-response systems became popular[14]. These systems build upon the whitelist approach. If a message is received from a sender not on

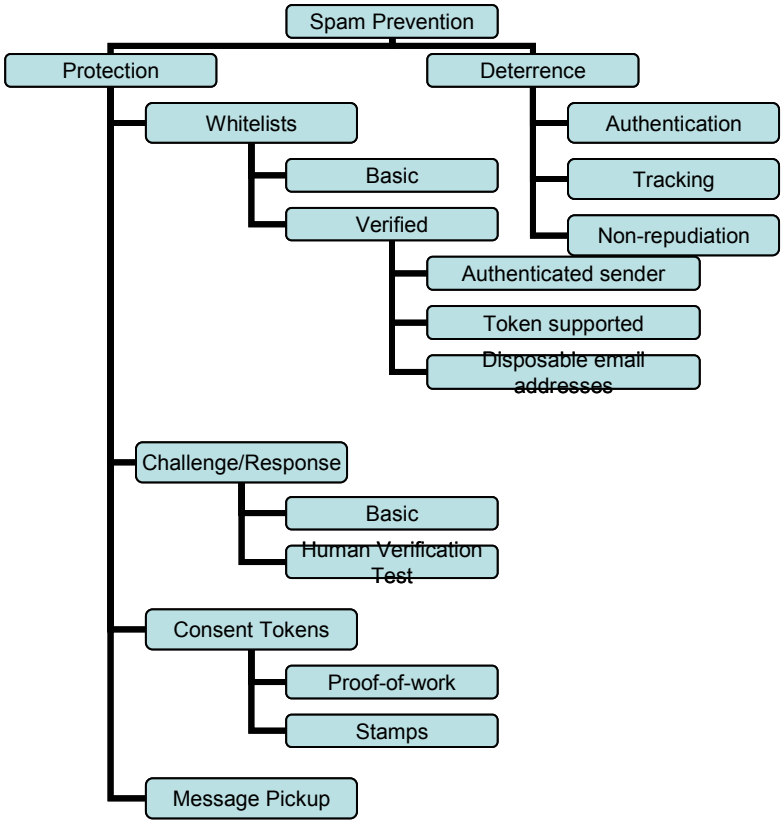

Figure 3: Types of Spam Prevention Systems

the whitelist, a 'challenge' is sent. A challenge is some message that requires a response in an attempt to verify the sender is human not simply spamming software. Consent tokens require that incoming messages arrive with some proof of effort such as a payment stamp or a proof of work. The goal of consent tokens is to require some effort (monetary, computational or otherwise) by the sender. Message pickup involves a system in which the emails are pushed to the recipient but instead the recipient picks them up.

Deterrence systems aim to discourage the act of spamming. This can be achieved by introducing authentication, email tracking, and non-repudiation[b] to the system. There has been significant activity placed on introducing email authentication into the messaging infrastructure. Suggestions for accomplishing this have gone as far to suggest a complete redesign or abandonment of email[15]. The more promising approaches focused on providing authentication on top of the current email systems rather than rebuilding email from scratch. There are two major paths being pursued in email authentication. To understand these, we first explain how email works without authentication in place.

In the email world, a domain publishes a list of machines that may receive mail on its behalf. This list is called the MX record and is the place messages are delivered for a given domain. There is no such authorized list of receiving machines for a domain. Therefore, any machine on the Internet can send emails claiming to be any other person or organization.

The first approach that we explore is domain-based IP authentication. This approach allows an organization to publish the list of authorized mail sending machines for a given domain. When an email is received claiming to be from example.com, the recipient queries the domain name service (DNS) servers to ask for the list of authorized senders for the domain. If the sending machine is not on that list, then there is likely a problem. Earlier versions of these systems came out of the Anti-Spam Research Group (ASRG) in 2003. In 2004, Microsoft introduced 'Caller ID for E-mail" then soon merged it with the existing Sender Policy Framework (SPF) domain authentication system to form Sender ID.[16]

The other leading approach for email authentication is cryptographic authentication of message headers. The major approaches to this were Yahoo's DomainKeys system and Cisco's Identified Internet Mail system. In July 2005, these approaches were combined to form DomainKeys Identified Mail (DKIM)[17]. DKIM allows an organization to publish a public key in DNS. Outgoing messages are signed using the organizations private key

[b] Non-repudiation is the property of guaranteeing that the sender of a message cannot later deny having sent the message. This can be achieved by using digital signatures or traffic trace back systems.

and the signature is inserted as a message header. The recipient is able to retrieve the organization's public key to verify the signature. This approach aims to achieve many of the same goals as IP authentication but has some strengths and weaknesses relative to that approach. The current best practice is that mail senders and recipients support both.

Early in the existence of email authentication systems, spammers were quicker to deploy the systems than legitimate systems, but today that trend is changing. For example, as of October 2004, 34% more spam domains published SPF records and passed the checks than good email. In April 2004, 11 of the Fortune 1000 had published SPF records. By October 2004, that number increased to 54 of the Fortune 1000. As of September 2005, 9% of the top-level domains had a Sender ID or SPF record and 0.5% had a DKIM or Domain Keys record. Email authentication has remained effective at identifying fraudulent email; For example, a message that fails a Sender ID check is ten times more likely to be spam than legitimate. [18]

3.3 Spam Response Systems

Spam response systems have grown to include several forms as shown in Figure 4. The traditional response to spam was deletion. In 2003, the use of other responses became widespread to improve effectiveness or perceived accuracy. Softer responses include labeling the message as spam or moving it to a junk mail folder. Today's systems allow a range of actions including quarantining the message for review by an administrator or end-user. This response grew in popularity in 2004 as the community became sensitive to false positives, or non-spam that is incorrectly detected as spam. As mentioned previously, challenge-response systems are increasingly used in some situations. Systems have been proposed to respond to spam messages by charging the sender.[19]

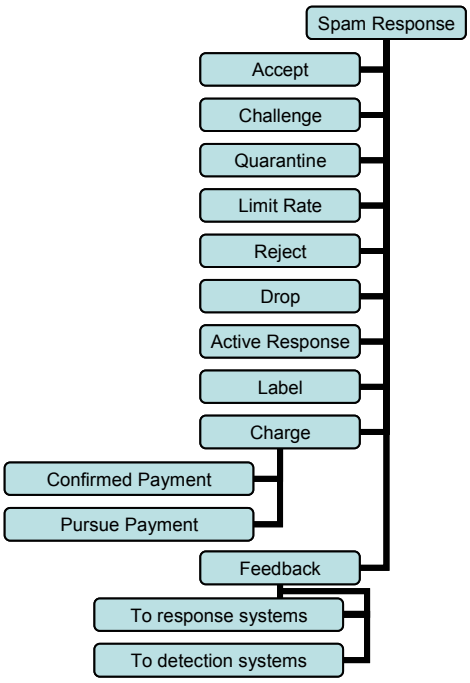

Figure 4: Types of Spam Response Systems

Active responses include traffic shaping or rate limiting to reduce the amount of traffic from suspicious senders. Traffic shaping can be used on ingress or inbound traffic to protect from external senders and on outbound or egress traffic to lock down internal computers. Traffic shaping refers to a set of techniques used to control the amount of traffic that a particular sender or set of senders can send in or out of a network. The amount may be measured in terms of factors such as number of connections, number of messages, or amount of data bandwidth. This is then enforced using some set of mechanisms such as slowing down the network connection from the sender, limiting the number of connections that will be accepted from the sender, or temporarily rejecting the sender's connection attempts. Ideally traffic shaping policies are only directed at senders that have been determined to be suspicious not across the entire traffic flow. Used

correctly, this technique can not only reduce the amount of spam that enters the network but reduce the amount of future spam that shows up at the organization.

The other use of traffic shaping or rate limiting is deploying at the egress point of an organization to control potential spammers within its network. Over the last few years this has become common among service providers and quickly led to scaling back some of the large spam sources located on their networks. The difficulty in this approach is tightening the rules enough to limit the usefulness for spammers while maintaining a usable system for legitimate senders. A related approach being adopted by service providers is the blocking of outbound port 25. This means that customer computers within their networks will not be able to connect directly to outside mail systems. Instead the customer machines will be required to send mail traffic through the ISPs mail servers to be delivered for them. This blocks zombie machines from being able to send messages directly out to the Internet.

3.4 Summary of Technology Changes

Overall the developments in spam detection, prevention, and response have advanced the state of the art since 2003. This has led to better protection for email users and created challenges for spammers. These technological improvements directly support CAN-SPAM and help make it more effective and vice versa. As both the legislation and technology push the spammers towards using a verifiable identity, it becomes easier to distinguish them from legitimate email senders. For example, one way the Act does this is by banning misleading headers and deceptive subjects. The technology does this with the use of authentication in spam prevention techniques. Similarly, just as the Act bans harvesting email addressing and relaying emails through other computers, spam response technologies have begun to respond more firmly to this activity.[c]

[c] For example, technologies are more focusing on directory harvesting attacks and messages originating from unknown or untrusted sources. As it has become more widely accepted that these activities are not allowed, technology is able to be more aggressive in dealing with them.

4 Trends Exhibited by Email Users

Beyond improvements in core anti-spam technologies, there are several broader factors that affect spam and anti-spam on the Internet. We call this set of parameters the email ecosystem. Here we will briefly examine the deployment of anti-spam systems, the use of mobile devices to access email, and the use of new messaging applications.

4.1 Deployment of Anti-Spam Systems by Users

Over the last two years the number of endusers, companies, universities, and ISPs that have deployed anti-spam systems has increased significantly.[20] This adoption of anti-spam systems has moved beyond large enterprises to small businesses. It is now considered a necessary cost of doing business on the Internet just as deployment of anti-virus and firewalls became several years ago. This wider deployment of anti-spam systems directly affects the success and profitability of spammers.

4.2 Use of Mobile Devices to Access Email

Since 2003, there is an increase in users accessing email via mobile devices. Many users of mobile devices still use traditional email clients as the primary means of interacting with email. There are differences in the message display and interaction capabilities on mobile devices versus traditional desktop email clients. These differences include the amount of the email subject that can be displayed, the ability to view message headers, and the ability to click web links.[d] Today there is no standard definition of these parameters since there are many different mobile devices and mobile email software used. Therefore it is difficult to give precise prescriptive advice to mail senders at this

[d] For example, in most devices the user can easily view the To, From and Subject, but it some devices it is not possible to view the extended headers such as X-headers and Received headers. This is the case for example on the Handspring Treo device using several different email clients. Also some devices have email access capabilities but not web access. This precludes these devices from accessing web links in emails.

point. However, senders should consider this when composing their messages. One possible best practice is that senders offer both an email based opt-out mechanism as a standard then supplement it with a web-based opt-out.

4.3 Other Messaging Applications Beyond Email

Email has become the most critical form of business communication today. Its uptime has become more important than even the telephone in many places. Spamming occurs in the email system because that is where the users are. As new messaging systems and protocols are developed and increase in popularity, it is a common cycle that the threats will appear in these systems as well. Already, a percentage of instant messaging traffic is spam. In the last year the number of instant messaging viruses has doubled. Other messaging systems include blogs, voice over the Internet, and wireless systems such as SMS. These other messaging systems are out of scope for this report because they are out of scope for CAN-SPAM.[e]

5 Trends Exhibited by Spam and Spammers

While from 2003 to 2005 the volume of spam sent on the Internet has increased, the amount of spam that reaches inboxes has decreased. The reasons for this decrease are improvements in anti-spam technology and increases in anti-spam deployments. The spammers have responded to the challenge of having their messages delivered in several ways:

1) Sending more spam: Since anti-spam systems are blocking a higher percentage of spam messages, from 2003 to 2005 spammers responded by simply continuing to send large volumes of spam. Simple multiplication suggests this will get more of their messages delivered to desktops. This largely explains the cause of the trend in increasing spam volume over the last few years. Spammers are not sending

[e] Some forms of VOIP and wireless SMS spam are covered by the Telephone Consumer Protection Act. Some SMS spam is within the scope of CAN-SPAM if it is sent to an email address.

more messages for no reason. They are doing so as an integral part of an effort to maintain or grow revenues.

2) Filter evasion: This refers to a wide array of techniques that spammers use to evade spam detection systems. This range from word misspellings, randomization of text, masking their identity.[21] It is evident from these techniques that the spammers employ skilled programmers that can understand the mathematics behind spam filtering algorithms and then develop countermeasures. This is the core of the 'cat-and-mouse game' of spam detection.[22]
3) Phishing: As we understand that spam is part of a business, phishing is an exercise in maximizing profits. Phishing attacks are more profitable than spamming because it increases the response rate and the revenue. The typical response rate for a spam campaign is a fraction of a percent while the response rate for a phishing attack is about 3%.

5.1 Spammers' Use of Zombies

Threat trends involve the new tactics, habits and tools of the spammers. One of the most significant shifts in the last two years is the use of zombie PCs to send the spam. Zombie PCs are innocent machine infected with a 'bot' program used for attacks. The 'bot' software typically reports to a controller channel and downloads and executes instructions from it. Once a zombie is created and controlled, it can be used by the attacker for several purposes including sending spam, sending viruses, creating mail relays, or harvesting passwords or email addresses. About 70% of all spam and 99% of all phishing is sent through zombies[23]

5.2 The Intent of Spam Messages

While CAN-SPAM focuses on commercial email messages, today unwanted messages are sent with a range of intentions. We refer to the traditional commercial emails as spam. Other unwanted messages are sometimes sent in bulk but are not commercial in nature. Phishing messages are those that pretend to be an organization with which the recipient

has a business relationship in order to extract information from the recipient. Malware is a somewhat general term often used to refer to any type of malicious code that is distributed.

5.3 Techniques Used by Spam Messages to Harm or Infect Computers

In this section we briefly explore the different methods used by spam to cause more damage or to infect computers. These include the user clicking a link within the message, using the unsubscribe link, or in some cases even if user simply views the message.

5.3.1 Harm Caused from Clicking Links in the Email

There is concern about the possibility of users infecting their computers by accessing opt-out links. Since the Act allows the sender to provide opt-out mechanisms as a web site, users may be required to click a web link in their attempt to unsubscribe. There have some examples in the wild of malicious content on the other end of the opt-out link.[24],[25] To get a better understanding of this issue, CipherTrust Research conducted an experiment that showed only 0.004% or about 1 in 25,000 spam messages contain opt-out links that lead to malicious content.[f] Since studies have shown that legitimate organizations respect opt-out requests, one suggestion is to have users opt-out from those senders while being more conservative about opting out from unknown senders. As new systems are being deployed that provide a graphical representation of the reputation information to the end user, users will have a better view of when to opt-out or not.

[f] The experiment analyzed 99,785 emails to extract 44,321 unsubscribe links. This produced 21794 URLS which led to 1719 downloadable files out of which 4 were found malicious. The 21794 links to the opt-out pages themselves are hosted on 4604 unique top-level domains and 13771 unique sub-domains. The 99,785 emails were from 81261 users representing 33038 email domains. To determine malign intent, we used a software behavior scanning engine that looks for inappropriate operating system calls and unsafe script behavior.

5.3.2 Harm Caused from Using the Unsubscribe Link

In this section we explore concerns that have been expressed about the risks of users unsubscribing to messages[26]. There are questions regarding the safety of opting out both by email or by web page. The first concern is whether opting out simply verifies email addresses and leads to the person ultimately receiving more spam in the future. There have been no studies to prove this theory and there have been studies that do a reasonable job of disproving this.[27] We suggest that due to the efficiencies of sending large numbers of messages from many zombie machines, verified addresses are not as valuable as they might have been in the past.

5.3.3 Harm Caused from Viewing the Message

Typically a worm is defined as a virus that does not require user action to infect the machine and propagate. There have been several examples on the Internet. For example, the Nimda worm exploited vulnerabilities in Microsoft Outlook to execute an attachment even if the user only previewed the message in the preview pane[28]. There are examples of this behavior that date back to at least the Kak worm in 1999 which used JavaScript and ActiveX to execute even if the user just read the message[29].

6 Conclusions

Over the last two years, technology advancements have improved spam catch rates and user experience while working well with the provisions of the Act. Spammer trends show that spammers have been affected by the Act and some are walking away from the spamming business while others face more difficult times in accomplishing their goal. Users have responded well by deploying anti-spam systems to defend their email systems and beginning to use best practices that will reduce their exposure to spam in the future. Moving forward users will continue to adapt and technology providers will continue to provide tools that fit well with the legislation to provide a robust combined approach to controlling the spam problem.

[1] M. Sarrel, "Lock Down Your E-Mail," PC Magazine, February 22, 2005.
[2] Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003, 15 U.S.C. § 7701 et seq.
[3] J. Postel,, "Simple Mail Transfer Protocol," RFC 821, August 1982.
[4] D. Crocker, "Standard for the Format of ARPA Internet Text Messages", RFC 822 August 1982.
[5] "The Email System and the Resulting Spam Problem," Federal Trade Commission, Appendix 2, Part III of the Commission's National Do Not Email Registry Report.
[6] D. Primack, VCs Get Piggy With It at Spam Smorgasbord**,** October 1, 2003, http://www.ventureconomics.com/vcj/protected/1060714633954.html.
[7] C. Metz, "Corporate Antispam Tools", PC Magazine, February 25, 2003.
[8] M. Sarrel, "Lock Down Your E-Mail," PC Magazine, February 22, 2005.
[9] Email Service Provider Coalition, "Project Lumos," September 2003.
[10] "Spam and Spammers", http://www.eff.org/deeplinks/archives/001774.php, August 2004.
[11] P. Graham, "A Plan for Spam," www.paulgraham.com/spam.html, August 2002.
[12] J. Zdziarski, "Ending Spam," No Starch Press, July 2005.
[14] L. Seltzer, "Challenge-Response Spam Blocking Challenges Patience", eWEEK, May 2003.
[15] Authenticated Mail Transfer Protocol, AMTP, http://amtp.bw.org, August 2003.
[16] "Sender ID", http://www.microsoft.com/mscorp/safety/technologies/senderid/default.mspx.
[17] E. Allman et al., "DomainKeys Identified Mail (DKIM)", Internet Draft, Work in Progress, October 23, 2005.
[18] "SPF Stats," http://www.ciphertrust.com/resources/statistics/spf_stats.php.
[19] B. Krishnamurthy, "SHRED: Spam Harassment Reduction via Economic Disincentives", http://www.research.att.com/~bala/papers/shred-ext.pdf, Working Paper, AT&T Research.
[20] B. Burke and R. Ryan, "Worldwide Secure content Management 2005-2009 Forecast update and 2004 Vendor Shares", IDC Market Analysis, November 2005.
[21] J. Graham-Cumming, "The Spammer's Compendium", http://www.jgc.org/tsc/, September 2005.
[22] Graham-Cumming, John, "The Spammers' Compendium", http://www.jgc.org/tsc/, September 2005.
[23] ZombieMeter, CipherTrust," http://www.ciphertrust.com/resources/statistics/zombie.php"
[24]E. Bangeman, "Can-Spam mandated opt-out link can lead to infestation" ars technica, "http://arstechnica.com/news.ars/post/20040922-4217.html", September 2004.
[25] J. Leyden, "Click here to become infected," The Register, "http://www.theregister.co.uk/2004/09/22/opt-out_exploit/," September 2004.
[26] Confidential 6(b) Order Response.
[27] "Top Etailers' Compliance With CAN-SPAM's Opt-Out Provisions", Federal Trade Commission, July 2005.
[28] W32/Nimda.gen@MM, McAfee, http://vil.nai.com/vil/content/v_99209.htm, May 2004.
[29] js/kak@m, McAfee, "http://us.mcafee.com/virusinfo/?id=description&virus_k=10509&affid=56," 1999.